

Home Network Technologies and Automating Demand Response

Charles McParland, Lawrence Berkeley National Laboratory

December 2008

DISCLAIMER

This document was prepared as an account of work sponsored by the United States Government. While this document is believed to contain correct information, neither the United States Government nor any agency thereof, nor The Regents of the University of California, nor any of their employees, makes any warranty, express or implied, or assumes any legal responsibility for the accuracy, completeness, or usefulness of any information, apparatus, product, or process disclosed, or represents that its use would not infringe privately owned rights. Reference herein to any specific commercial product, process, or service by its trade name, trademark, manufacturer, or otherwise, does not necessarily constitute or imply its endorsement, recommendation, or favoring by the United States Government or any agency thereof, or The Regents of the University of California. The views and opinions of authors expressed herein do not necessarily state or reflect those of the United States Government or any agency thereof or The Regents of the University of California.

Home Network Technologies and Automating Demand Response

Charles McParland, LBNL

Commercial and Residential Demand Response

Overview of the Demand Response Landscape

Over the past several years, interest in large-scale control of peak energy demand and total consumption has increased. While motivated by a number of factors, this interest has primarily been spurred on the demand side by the increasing cost of energy and, on the supply side by the limited ability of utilities to build sufficient electricity generation capacity to meet unrestrained future demand. To address peak electricity use Demand Response (DR) systems are being proposed to motivate reductions in electricity use through the use of price incentives. DR systems are also be design to shift or curtail energy demand at critical times when the generation, transmission, and distribution systems (i.e. the “grid”) are threatened with instabilities. To be effectively deployed on a large-scale, these proposed DR systems need to be automated. Automation will require robust and efficient data communications infrastructures across geographically dispersed markets. The present availability of widespread Internet connectivity and inexpensive, reliable computing hardware combined with the growing confidence in the capabilities of distributed, application-level communications protocols suggests that now is the time for designing and deploying practical systems.

Centralized computer systems that are capable of providing continuous signals to automate customers’ reduction of power demand, are known as Demand Response Automation Servers (DRAS) [1]. The deployment of prototype DRAS systems has already begun – with most initial deployments targeting large commercial and industrial (“C & I”) customers. An examination of the current overall energy consumption by economic sector shows that the C & I market is responsible for roughly half of all energy consumption in the US. On a per customer basis, large C & I customers clearly have the most to offer – and to gain – by participating in DR programs to reduce peak demand. And, by concentrating on a small number of relatively sophisticated energy consumers, it has been possible to improve the DR “state of the art” with a manageable commitment of technical resources on both the utility and consumer side.

Table 1. US Energy Consumption by Sector (2009 – 2010) [2]

Sector	Percent Energy Use
Residential	22%
Commercial	19%
Industry	31%
Transportation	28%

Although numerous C & I DR applications of a DRAS infrastructure are still in either prototype or early production phases, these early attempts at automating DR have been notably successful for both utilities and C & I customers. Several factors have strongly contributed to this success and will be discussed below.[3] These successes have motivated utilities and regulators to look closely at how DR programs can be expanded to encompass the remaining (roughly) half of the state’s energy load – the light commercial and, in numerical terms, the more important residential customer market. This survey examines technical issues facing the implementation of automated DR in the residential environment. In particular, we will look at the potential role of home automation networks in implementing wide-scale DR systems that communicate directly to individual residences.

Commercial DR Protocols

As mentioned above, given the level of peak demand savings achievable for large C & I customers, commercial DR programs have received the greatest initial attention. One important feature in this domain has been the development of the OpenADR (Open Automated Demand Response) communication’s standard [4]. This application-level communications standard, based on IP protocols, has been under development over the past six years and has, in large part, served as an organizational forum for promoting interoperability among existing and new commercial DR technologies and existing control systems in commercial buildings. OpenADR defines and codifies the messages exchanged between an energy supplier and participating consumer’s during periods of desired demand reduction. By building on well established information technology (IT) and communications standards, in particular, the XML markup language and the Internet protocols, this effort has been able to rapidly progress from a proposed standard to an operational one. It presently represents the most fully developed infrastructure for both the role and the requirements of network communications in creating

operational, distributed DR systems. The following list highlights the major characteristics of the OpenADR application level protocol.

- **Continuous and Reliable** - Provides continuous, secure, and reliable two-way information exchange infrastructure in which the clients at the end-use site receive and return acknowledgement to the DRAS (DR Automation Servers) upon receiving DR signals.
- **Translation** - Translates DR price and event information into continuous 5-minute Internet signals to facilitate automation of DR. These signals are designed to interoperate with facility Energy Management and Control Systems (EMCS), lighting controls, or other end-use control devices.
- **Automation** - Receipt of an external signal is designed to initiate automation through the use of pre-programmed DR strategies determined and controlled by the end-use customer.
- **Opt-Out** - Provides opt-out or override function to consumers if the price or DR event comes at a time when reduction in end-use services is not desirable.
- **Complete Data Model** - Describes a rich data model and architecture to communicate price, reliability, and other DR activation signals.
- **Scalable** - Provides an architecture scalable to different forms of DR programs and tariffs.
- **Extensible XML-based messaging** – Standardize server/client message formats using the XML markup language and select XML tags to allow future semantic and content extensions to OpenADR messages.
- **Open Standards** – Utilize open standards-based technology such as Simple Object Access Protocol (SOAP) and web services form the basis of the communication standard.

The advanced state of this DR application protocol has important implications for the utility operations and/or aggregator side of the DR equation. At a systems level, it is clearly advantageous for all potential DR customers or clients to appear logically identical. This greatly reduces the complexity of treating each customer in a unique manner that requires specialized software. As OpenADR achieves practical success in the large C & I world, it can be argued that when expanding into the light commercial and residential domain, DR customers in that domain should look, at least from the utilities perspective, the same as existing C & I ones. While this may well be a reasonable design goal on the utility or DR server side, we will show that there are a number of differences on the consumer or DR client side that require close examination. A discussion of the differences between the large C & I and residential customer environments will help illuminate this issue.

Residential vs. C & I DR Environment

As noted above, there is growing interest in applying the principles and, to the extent possible, the technology developed for large C & I customers to the residential market [5]. Residential demand reduction programs have been around, in several forms, for a number of years. For example, a number of utilities and aggregators have implemented direct-load control to allow utility-initiated control of residential loads and provide real-time monitoring of the requested load shedding behavior. While these programs do provide utilities and aggregators with known and contractually agreed upon demand reduction capabilities, they have not achieved large participation in the residential market. It is not unusual for these load control systems to be de-energized or bypassed at the appliance for reasons of convenience or necessity [6]. And, since they are, for the most part, not the responsibility of the end user, failures or maintenance issues receive little attention.

There have also been third-party efforts to enroll households in DR programs that promote, through custom instrumentation and controllable appliances, reduced household energy and peak demand use. The resulting aggregate energy savings of all participating households is then sold back to the utility at a pre-agreed price per unused kW. Again, while there has been some success with these programs, the majority of the residential market remains unenthusiastic about joining these programs and having customized equipment installed in their homes. Overall, there is general agreement among regulators that less intrusive and more flexible DR mechanisms will be necessary to insure truly wide-scale residential participation in a DR program [7]. In part, this survey is a beginning step towards

understanding how existing and future home automation systems may provide the necessary, user-accepted infrastructure required for truly wide-scale residential DR systems.

At the highest level, DR systems for both C & I and residential areas share many of the same operational goals. However, there are a number of important differences between these environments and the successful implementation of residential DR systems will depend on a close examination of these differences. For one thing, the C & I and residential domains bring different resources to the task of managing energy use. Energy costs in owner-occupied commercial buildings are often evaluated by sophisticated energy managers. So, in principle, business mechanisms already exist for regularly reviewing these costs and, given appropriate tools, management staff can be tasked to evaluate the value and effectiveness of a DR implementation and judge its effect on occupant comfort. In the residential environment, these costs are typically only evaluated when utility bills arrive and, given very nature of everyday home life, household decisions to alter energy consumption patterns often lack the clarity and level of commitment found in the commercial world. Furthermore, the considerations used when evaluating the residential “comfort” gained for cost expended metric are often unclear and easily reversed. When installing DR equipment in large commercial buildings, resident facility staff is typically available to manage the control company’s efforts to implement DR-required changes in a building’s energy management control system (EMCS). In residential settings, technical expertise is often lacking and the usual industry metrics for evaluating energy usage are often felt inappropriate or clinical (e.g. OSHA workplace environmental standards). As a result, residential DR operations are not approached with the same diligence found in commercial settings and the full potential energy (and dollar) savings may not be achieved or even recognized. While the underlying fundamentals may be identical, residential DR implementations will have a greater need for transparent and user-friendly operation than those found in commercial buildings.

Another difference between the residential and commercial DR domains is the dramatic variation between the potential number of participants in each. As noted earlier, the number of residential customers will exceed the number of large commercial customers by at roughly a factor of 40. As commercial DR implementations, thanks to developments like OpenADR, become more prevalent, the number of DR participants will scale with the number of commercial and industrial buildings in a utility’s service area. It is likely that, initially, new participants will come from large, multi-tenant buildings whose current staff can manage the additional effort needed in configure their building. Commercial buildings with sub-metering will be very strongly motivated to participate while others may delay until the financial benefits of DR can be allocated to individual tenants. So, initially, the aggregate number of participants

will be less than the total number of large C & I customers in a given service area. Residential DR participants, however, will scale with the number of homes in a utility's service area and that number is noticeably larger than that of commercial buildings (about 1.8 million one- and two-unit dwellings in 9 county Bay Area alone) [8]. If utilities and aggregators succeed in making residential DR attractive to residential customers, aggregate numbers of DR clients can grow very quickly (we need to discuss this sentence). The differences in scale between wide-spread deployments of DR systems to the commercial vs. the residential market will have a large impact on IT and communications infrastructure costs associated with their respective operation. A successful residential DR architecture will need to carefully address the cost and reliability issues associated with these, potentially, very large client bases.

It has been proposed that, at least in the residential domain, actual DR clients be aggregated in some fashion and appear to utilities as a smaller number of OpenADR clients. Such architectural accommodations will prove beneficial by reducing the number of individual clients requiring constant communication. It should be noted that, in the current OpenADR design, each DR client is periodically probed by the DR server to verify its presence. So, if multiple residential DR clients are somehow represented or proxied by a single OpenADR client, then overall number of periodic network transactions initiated by the OpenADR DRAS system will be reduced. However, to some extent, this approach simply pushes the problem downstream. If the semantics of OpenADR operations requires a continuous indication of which participating OpenADR clients are on-line and potentially responsive to DR events, then additional mechanisms will be required to collect this information and make it available to the DR server. This may include polling the OpenADR proxy of its clients through its private DR event communications channel or, perhaps, contacting DR controllable end points through other signaling mechanisms (e.g. Internet). These mechanisms do not, at present, exist and will need to be designed and, since they ultimately reflect the availability of potential DR responders, reconciled with the semantics of OpenADR. In short, while there may be technical approaches that make it more feasible to marry commercial OpenADR systems to residential DR applications through some form of aggregation, there may be difficulties in representing the DR "client" in a consistent and transparent manner across both commercial and residential domains.

The inherent scaling issues of residential DR also surface in the area of data collection. While the current OpenADR standard does not address real-time utility collection of data from DR clients, the communication standard on which it's based, the Internet TCP/IP protocols, will support such capabilities. If, in the evolution of commercial DR systems, it becomes necessary to feed customer DR responses and current load information back to utilities, such data

collection can, given the existing Internet communications infrastructure, be accomplished in a number of ways. The residential domain, however, offers a different set of challenges. While the aggregation of DR clients behind OpenADR proxies described above may allow a smooth deployment of DR into the large residential domain, the potential of adding real-time data collection to the residential DR picture raises several design concerns. If there is a genuine utility requirement for near real-time monitoring of end user loads, then residential DR will also require two-way communications channels between utilities and those loads. This requirement will greatly reduce or eliminate the potential for one-way radio transmission as a medium for DR event dissemination. As a result, a system level requirement for detailed collection of end residential load data effectively removes a low-cost, potentially ubiquitous communications channel from consideration for DR system deployment. This further illustrates the coupling between DR system design and system scaling in residential DR applications. If, as some have suggested, residential load monitoring need only be examined at aggregate (e.g. substation or feeder) levels, then such restrictions can be removed and the system level design can be freed from this communications technology constraint.

Home Networking Systems

What are Home Networks?

The term “home network” refers to a number of data communication and control technologies that have been developed primarily to serve applications in the home. Over several decades, these systems have evolved to provide increased convenience to home occupants in at least three major areas: simplified remote control of residential entertainment components and systems (TV, home theater, audio, etc.), automated remote control of major house sub-systems (thermostat settings, garden irrigation, security, etc.) and, most recently, flexible Internet connectivity to IT devices throughout the home (PC’s, printers, etc.) and the greater Internet. While the terms “home networks” and “home automation” are frequently used interchangeably, this paper will attempt to differentiate them in the following manner. Home automation will refer to a system, product or application that has been designed to control devices within the home. Home networks will be used in a much narrower context and refer, primarily, to the particular physical and logical signaling used to convey information between elements within a home automation system.

Whenever data and control networks are discussed, it is inevitable that the network being examined is compared to the standard OSI seven-layer network model. This model is a well-accepted conceptualization of all the essential tasks required of an ideal network. As an

abstract model, the OSI model provides a common way of organizing and separating the software and hardware network components into manageable, compartmentalized layers. [9] While designers can choose to simplify or combine tasks performed by these layers based on the requirements of the network that is being created, the strength and usefulness of the OSI model stems from its completeness – there is a layer for every major network function. If layers are left unimplemented or improperly combined, it will be very difficult to reliably interoperate with other networks that have fully designed in compliance with the OSI model.

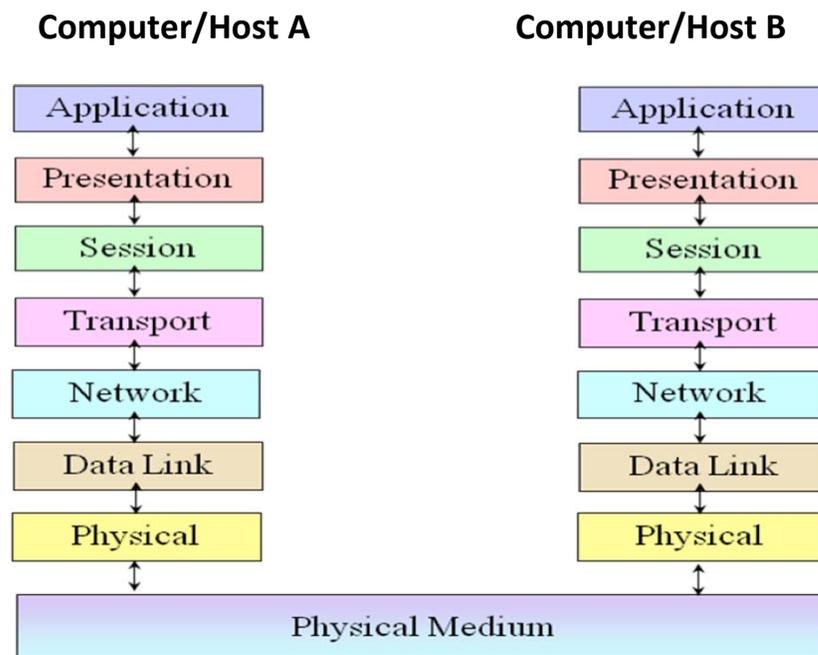


Figure 1. OSI 7 Layer Reference Model

While the OSI model is well-accepted as the most complete conceptualization of an “ideal” network, many network protocols do not fully implement it. For example, implementations of the very successful Internet protocol, TCP/IP, typically combine functions of several OSI model layers into a single TCP/IP layer. Since experience gained in the development of TCP/IP substantially influenced the design of the OSI model, these differences are not considered major and the Internet protocols are, for practical purposes, well-aligned with the OSI model. Unfortunately, many home networks and control systems have historically been designed with little or no regard for these functional layers. Typical home automation systems have been conceived more as “remote control” applications than as actual data networks.

Although they operate by communicating digital data between computers (e.g. microprocessors) and are, in a very real sense, actual data networks, these systems were never thought of as such and, as a result, were designed without any reference to the above OSI network model. For example, a home automation system that provides remote control of a thermostat set point will appear very unlike the network running on a personal computer (PC) and more like the following:

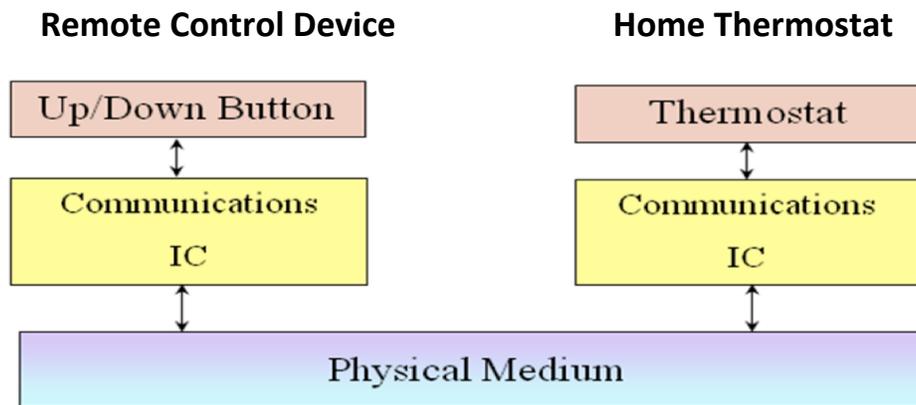


Figure 2. Non OSI Model Network Layering Example

The functional behavior of each of the above layers is solely defined by the “application” the device is implementing with no real consideration of data network issues. A different home automation system, for example, a remote pool pump switch, would have a network layering diagram similar to the above. But, lacking any effort to implement network layers that describe how data should be routed through the network and how it should be presented at each end, it would be difficult to combine these two devices into a single home-wide control system. The difficulty of combining or networking different devices that do not share the same conceptual layering is clear, and a comparison of the above diagram to the OSI 7-layer model reveals the difficulties that would be encountered in combining such devices with other modern OSI-compliant devices such as a PC.

For cost and marketing reasons, the home automation domain has many such non OSI-compliant devices that will prove resistant to efforts to incorporate them into modern, data networks. Fortunately, as new network protocols are developed and become standardized, that situation is changing. As will be shown below, recent initiatives to design low-cost networking devices that are suitable for home automation applications are using modern, rigorous networking models similar to the OSI model. However, in considering the present state of home networking, the network differentiates between single-purpose, proprietary

devices. Modern OSI-like protocols will prove to be a critical issue in understanding how home automation systems can and cannot evolve and this issue should be kept in mind for the remainder of this paper.

Evolution of Home Networks

While industrial-grade control networks are often required for safe operation of large C & I facilities, home networks are not, at present, “needed” for single-family residences to operate as comfortable, habitable living spaces. The value proposition for home networks is primarily based on convenience to occupants. Many areas of home automation are functionally synonymous with the “remote control” and these house-wide remote control systems are, in fact, rudimentary networks that function by transferring data between controls and appliances within the house. And, while “convenience” may be a difficult metric to evaluate, it has driven the development of a burgeoning group of manufacturing, installation and sales businesses. The pursuit of convenience at an acceptable cost has resulted in household products that address a surprising variety of home operations. However, only some of these operations have great relevance to potential DR systems and reveals yet another difference between the commercial and residential DR domain. At present, home networks support automation systems whose primary motivation is convenience. And, while these same systems will prove beneficial in addressing residential DR design goals, the original intent for their deployment to the home, namely convenience, should not be forgotten or subverted.

Since the convenience metric is difficult to compare by cost, it is not surprising that, in attempting to optimize the “cost of convenience”, the home automation field has leveraged a wide variety of inexpensive communications and networking technologies to support each product’s particular needs—low bandwidth protocols for simple device control (e.g. on/off, volume, etc.), high bandwidth protocols for delivery of audio or video to multiple rooms, etc. Unfortunately, these network technologies have not evolved from a common set of root standards. In fact, many of them were developed and marketed at a time when the benefits of open standards were not yet fully appreciated. As a result, most home networking protocols exist completely outside of any standards sanctioning body. Marketing and licensing issues have limited the growth of many of these network protocols and, at least at the simple functionality end of the spectrum, it is not unusual to find only a single manufacturer responsible for providing the electronics required for media access (MAC) and physical access (PHY) functions.

This stands in stark contrast to the third category of home networks – one that has not been designed specifically for residential use but is increasingly found in many homes. These networks are the family of open standard (logical and physical) Internet protocols (IP) that are at heart of every PC, laptop, printer and router found in the home or office. The success of these networks in the marketplace has been driven primarily by the growth in personal computing and has been truly remarkable [10]. And, with their explosive growth and increasing affordability, they have migrated into households in increasing numbers. While brought into the home primarily as the backbone for a residential “IT” (information technology) networks, they also represent the most fully-engineered, robust and standardized network technology that can be found in the home at present. As these protocols become common in low-level industrial control applications and migrate into audio/video appliances, few will argue against their ability to serve all the needs presently found in home automation applications. But, for reasons of cost and/or complexity, these IP-based networks remain largely on the sideline for most home control applications.

While the popular wired IP/Ethernet protocol standards of the IT world may have proven too costly for low-end home networking applications, a new breed of wireless, low-power networking devices has emerged from the development labs and are finding their way into residential networks. These network devices, based on research aimed at developing low-power, wireless sensor networks (WSN), have been through the open standards process and, at least for the lowest protocol layers, conform to a publicly available specification (IEEE 802.15.4) [11]. Both individual MAC/PHY (i.e. data link and physical layer) chips and assembled modules have been shipping from multiple manufacturers for several years. Since these devices have targeted low-cost applications, it is hoped that they may well find acceptance in low-end, residential control applications. However, while the lowest layers of this protocol are open standard, the higher layers that, among other things, control message formatting and presentation of data to users or computers remain non-standard. A number of major efforts are underway to standardize these upper protocol layers and, if successful, a modern, inexpensive, and interoperable low-end network protocol could quickly become the basis of most home automation systems.

Intersection of Demand Response and Home Networks

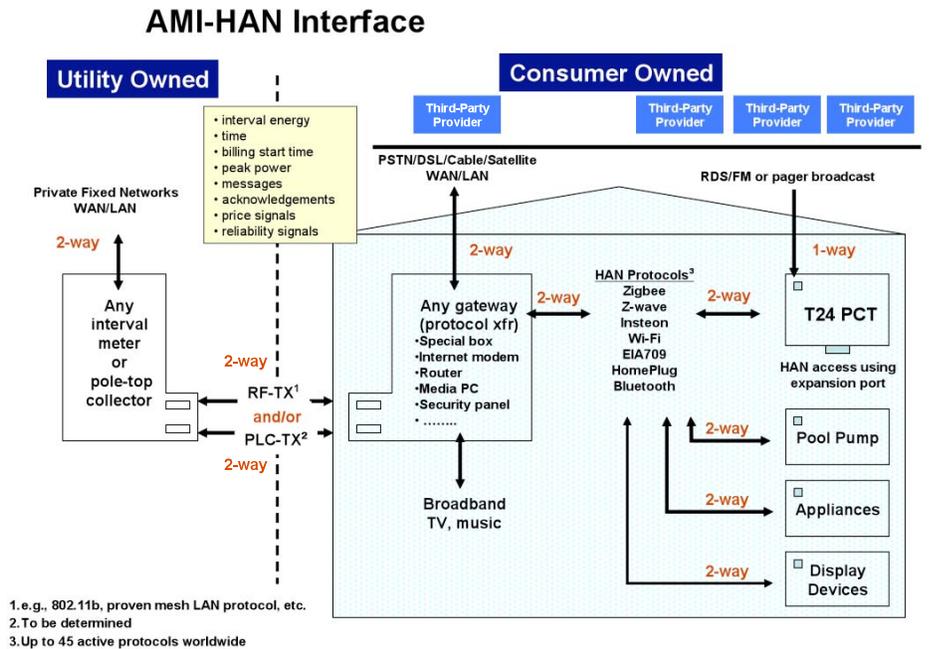
Last “Hundred Feet” Communication Problem

As a core part of their business, utilities and allied service providers have developed the infrastructure needed to distribute energy, water and gas to their residential customers. Over

the course of many decades, a reasonable business and liability model has evolved that describes where a utility's infrastructure ends and the residential distribution system begins. The result is a utility "service termination point", usually located on the residential periphery, where consumption is metered for revenue purposes. Distribution of utility consumables (electricity, gas, water, etc.) beyond that point to the actual site of consumption is primarily the responsibility of the home owner and not the utility.

However, the desire to use utility-generated DR signals to modulate power consumption implies that control signals, or their logical equivalent, will be propagated all the way from the control requestor (e.g. utility operations office) to the point at which energy consumption is controlled (e.g. appliance power switch). In other words, air conditioners and other major appliances within the home need to receive and properly interpret signals that are requests from utilities, or their surrogates, for demand reduction. Exactly how this will happen and who will have responsibility for each part of the path through which control signals pass will greatly effect the system's deployment and operation. While this point may appear obvious, what is not clear, at present, is who will be responsible for and, ultimately control the communications path ("the last hundred feet") within the residential envelope that conveys these signals to the energy load. Unlike the utility meter that marks the boundary of a utility's responsibility for power distribution, there is no accepted model that describes where responsibility for control signals change hands.

Utilities and researchers have been investigating several communications technologies capable of passing signals into the "residential envelope". These include a variety of techniques from message transmission by broadcast (e.g., Radio Broadcast Data System or "RBDS") [12] and by wideband via the Internet to local narrow-band power line communications protocols (e.g., HomePlug) and low power wireless protocols (802.15.4) emanating into the "residential envelope" from the revenue meters located at the utility/residence termination point.



Ron Hofmann

Figure 3. AMI-HAN Interface

As technology enables the communication of remote control signals into the residential envelope, the issue of “who is responsible for control signals in the home” will become increasingly important and, furthermore, the question of who can switch loads and whether that control is shared will need to be clearly addressed. While it is clear that direct and exclusive utility control of device loads within the residential envelope may create a control system with predictable DR behavior, this control model will create a residential environment at odds with the expected convenience-driven behavior that has come to be synonymous with home networks and automation. Overall, it appears that, if the responsibility over power control signals resides with external agencies while power distribution infrastructure remains the responsibility of the home owner, confusion, dissatisfaction and erratic system behavior will be certain.

It should be pointed out, that this is the same logical space now being pursued by the growing home automation effort. Although currently motivated primarily by convenience, home network protocols have been engineered to provide cost effective, short distance communications capable of controlling home loads from various locations in the home. The existence such a manufacturing and marketing base holds great promise for implementing residential DR systems. Logically, since these systems are capable of communicating user

command from various locations in the home to remote loads, they should be readily adaptable to accepting inputs from utility interfaces located at the edge of the residential envelope. But, at a systems level, there still remains ample opportunity for serious collisions between commands emanating from multiple sources. Central to this issue is whether the home is viewed as a single control nexus that receives, interprets and distributes utility DR signals within the residential envelope or whether individual controllable loads within the home are viewed as utility control targets undifferentiated with respect to the particular home in which they are located. If competition for “appliance control” is allowed within the residence and entities inside and outside the home vie for control over multiple “last hundred feet” communications media, then who or what will arbitrate between conflicting commands?

Examination of Home Networking Technologies

Evaluation and Selection Methodology

As noted above, home network technologies have evolved to service a wide variety of applications. By focusing on network technologies that either already appear in the home or are being considered as candidates for home deployments, we have essentially selected technologies whose capabilities and manufacturing costs are suitable for residential applications. So, while not explicitly using unit or system costs as criteria for inclusion, the technologies examined here reflect those most likely, at present, to be integrated into residential DR systems.

It should be pointed out that one particular cost characteristic has already greatly influenced design and innovation in the domain of home networking – the cost of system installation. Historically, residential design and construction techniques have been slow to adopt building practices that promote easy and inexpensive installation or retrofit of wired sensors and controllers throughout the home. Although this shortcoming is starting to be addressed by the use of structured cable bundles (e.g., combined video, computer networking, analog phone, etc.)[13] in new high-end home construction and major residential remodeling efforts, most homes lag far behind commercial structures in terms of accessible utility and electrical wiring corridors. As a result, many designers of home network systems have rejected dedicated-wire communications technologies as impractical and, ultimately, too expensive to foster large-scale acceptance across existing and new housing stock. While a number of reliable, well-established networking technologies, including many hard-wired industrial control standards, are ubiquitous in commercial and industrial settings, they are considered impractical for whole house applications. Although wired computer data communications standards, such

as those used in wired Ethernet (10/100/1000BT), continues to grow, in our opinion, the field of home networking will be dominated by standards that focus on minimizing installation costs – in short, wireless or power line communications technologies.

Discussion of Network Characteristics

Network terminology can, on occasion, be at odds with usage in other domains. In the field of data communications, a wireless network is one that utilizes some form of modulated RF to pass data between communicating nodes. While it is also possible to “wirelessly” communicate with infrared light (e.g. most consumer appliance remote controls), line of sight interference makes communications subject to frequent drop outs and, therefore, infrared signaling is seldom used. As a result, in data communications and networking domain, wireless networks are assumed to be based on RF communications technologies.

This can be contrasted with a different usage of the term “wireless” – namely any communications technology that requires “no new wires”. These networks include technologies that allow multiplexing additional information onto existing wire or cable infrastructure. In almost all cases, the design goal for a “no new wires” communications system is to add additional functionality to devices that are already electrically connected. This technique has been adopted by home entertainment networks to distribute remote control signals (low frequency control data, approx. 1 kHz.) over the same coaxial cable used to distribute television RF (radio frequency) signals (54 – 1002 MHz.). And the widest, and most creative, use of this technique can be seen in the myriad of standards used to add data communications to existing AC power distribution wiring – “power line carrier” protocols. The presence of large loads switched by bare contacts typically creates challenging interference spikes that cause data reliability problems for most of these data communications protocols and, as a result, there are multiple open and proprietary standards that attempt to mitigate these issues in a variety of ways. For the purposes of this survey, we will refer to these as “power line protocols” and not refer to them as wireless (i.e. no new wires).

The following terms will appear in evaluation tables. Detailed discussion of table contents will follow each section.

Bandwidth: In broad terms, bandwidth attempts to quantify how much data can be transferred over a communications channel in a given interval. Some of the bandwidth for a given channel is used to support the particular data protocol being used. This usually includes data for addressing or routing information as well as cyclic redundancy check (CRC) and checksum fields that allow verification of the correctness of received data. Most residential load control

applications require relatively little communications bandwidth. Historically, systems requiring dedicated closed-loop control, have been designed to function without the need for home networks. Examples include basic heating and air conditioning systems and hot water tanks. And while home networks are capable of interfacing with these systems to alter operating set points, there is no requirement to insert home networks into the primary control feedback loop for these systems. Within the home network domain, the only systems requiring high network bandwidth are those associated with multimedia (audio and video) and, within the current DR discussion, they have little relevance.

One-way, two-way: In an attempt to minimize cost and complexity, some home networks have implemented systems that only transmit data in one direction – from a transmitting controller to a receiving controller. The result, at a system level, is that while it is possible to issue a network command to turn a load on or off, the network does not have the capability to determine whether that command was correctly received and, if so, correctly acted upon. In some applications, it is acceptable to have load control “mostly work” (e.g. switching on room lights). But, it is not clear if DR applications will accept this low level of robustness.

Security: This is a complex issue that, given the wide variation in functionality found in home networks, does not cleanly cut across all candidate technologies. There are essentially three issues. First, can the receiving node verify that data received is exactly the same as the data packet sent by the transmitting node? Second, can the receiving node be certain that, in receiving data that has been spread over multiple logical messages, no intervening messages have been lost. And third, can the sending and receiving node package message in such a way that the designated receiver can verify that the message only originated from a trusted source and that only the designated receiver can, in fact, read it. Many of these characteristics are missing from simple home network protocols. Some protocols are adding security features. And, of course, high-level, IP-based protocols are capable of implementing all of these features. While security may not be considered a critical feature in controlling residential loads, it has serious implications for wireless technologies that might allow adjacent dwellings or outside computers to penetrate the “residential envelope” and control critical house functions without the resident’s knowledge. (We need to distinguish security “to” the home from security “in” the home.)

Openness: This is an indication of the public or private nature of the network protocol specification. The primary implication is that open protocols are more likely to engender widespread adoption and richer set of products in the marketplace.

Protocol Specification status and Completeness: Is the protocol available in written form or is it only standardized by example designs produced by the license holder? Is the protocol mature enough to support commercial products?

Developer Protocol Interface: This characteristic primarily addresses how a system designer gains access to the protocol. Is it possible to create entire protocol messages within a user interface PC or display device and then send them to a target node, or do I need an application level command and, by some mechanism, pass that command to a manufacturer's chip that subsequently formats and send an appropriate message?

Deployment and Diagnostic Tools: Are deployment and diagnostic tools available? Since, especially with wireless technologies, home network robustness is an important concern, diagnostic tools are critical to simple, efficient network installation.

Application Programmers Interface (API): In general, even the simplest of home automation system designs can become quite complex when they are actually implemented. It has become common practice in well-designed software systems to create carefully designed interfaces at several key locations in the logical program "flow" that spans the distance between the user interface and the lowest level of interaction with hardware. These software interfaces define a set of capabilities and establish the semantics for how the system behaves when each interface is called. While these stable places appear between all layers of a complex software design (e.g. the OSI seven layer model), they were originally designed to facilitate and stabilize the partitioning of complete software packages into an application and system layer. The phrase "application programmers interface" or API has therefore become synonymous with these important code separation boundaries.

Wireless Protocols

In general, the wireless protocols shown below fall into two groups. One group owes its lineage to IEEE Ethernet standards that evolved to include various wireless media formats. The other group evolved from proprietary developments of wireless “remote control” technologies that, for cost and marketing reasons, did not address standard data network design issues. As noted above, complete integration of these two classes of protocols into a single robust system will be difficult and expensive. However, devices found in the second group of protocols can probably function as simple logical links to remote sensors and actuators. Thus, a reliable, home-wide automation system could be designed by judiciously combining these two protocol groups and limiting the degree to which proprietary protocols penetrated into the overall design.

Although not elaborated in the following tables, IEEE 802.15.4 plays a fundamental role in several of the following protocol stacks. Thus, at the PHY level, signals in the 2.4 GHz. RF spectrum are governed by 802.11 Wi-Fi [14], 802.15.1-3 (Bluetooth) [15] and 802.15.4 [16] low level standards. The upper layers of each of these protocols are tailored to the specific design goals of each standardization effort. So, even in the case of protocols that share the same PHY and MAC layer (e.g. ZigBee and 6LowWPAN), interoperability between these protocols at the application layer is, at best, extremely difficult and, in practice, unattainable.

Also, it is worth noting that some standards, such as WiFi and 6LowWPAN stop at the definition of socket-like abstractions when defining the application API. Others go further and define the semantics and exact programming syntax to be used when an application interacts with the underlying protocol layers. While the existence of these usage “profiles” can speed the development of individual products, they may also be seen as overly proscriptive of the application design. The style and scope of the application interface defined by each of these standards is a primary differentiator when matching these protocols to application/product requirements.

Table 2. OSI layers vs. Protocol Standards

OSI Network Layer	Proprietary ZigBee	IETF 6LowWPAN	IEEE-802.15.1-3 Bluetooth	IEEE 802.11 Wi-Fi	Proprietary Z-Wave
Application	Application / Profile API included in standard	Application API included in standard	Application / Profile API included in standard	IP	Application / Profile API included in standard
Presentation	Part of standard	Part of standard	Part of standard	IP	Part of standard
Session	Part of standard	Part of standard	Part of standard	IP	Part of standard
Transport	Part of standard	Part of standard	Part of standard	TCP/IP	Part of standard
Network	Part of standard	Part of standard	Part of standard	TCP/IP	Part of standard
DataLink (MAC)	IEEE 802.15.4	IEEE 802.15.4	Hardware/firmware	Hardware/firmware	Part of standard
Physical (PHY)	IEEE 802.15.4	IEEE 802.15.4	Hardware/firmware	Hardware/firmware	Part of standard

Table 3. High-Level Protocol Characteristics

Protocol	Bandwidth	One/Two Way Communication	Security -Data Integrity -Protection from Data Loss -Confidentiality	Openness
ZigBee	250Kbps max.	Two way	-predicted to be very robust -security and data loss status unknown - application templates	Private, available to license holders
6lowpan	250Kbps max.	Two way	-Full IP security capabilities	IETF standard
Bluetooth	1-3 mbps	Two way	-Full IP security capabilities	IEEE standard
Wi-Fi	2 to 248 mbps	Two way	-Full IP security capabilities	IEEE standard
Z-wave	9.6 or 40Kbps	Two way	-Minimal security addins in progress	Private, available to license holders

Table 4. Protocol Openness and Usability

Protocol	Protocol Specification Status	Protocol Status/ Completeness	Developer Protocol Interface	Deployment / Diagnostic Tools
ZigBee	Written spec. available to licensees	Early chip level implementations shipping	802.15.4 plus assembled modules, multiple vendors	Yes
6lowpan	Written spec.	Complete with routing protocols in progress	802.15.4 plus assembled preliminary modules, multiple vendors	Yes
Bluetooth	Written spec.	Complete	MAC/PHY chips and modules from multiple vendors.	Yes
Wi-Fi	Written spec.	Complete	MAC/PHY chips and modules from multiple vendors.	Yes
Z-Wave	Written spec. available to licensees	Complete, security additions in progress	Microcontroller interface available under license	Yes

Power Line Protocols

The prior comments for the wireless protocols apply to power line protocols as well. Interestingly, there is a greater predominance of proprietary protocols in the power line communications area (see tables below). Historically, it has been difficult to engineer power line protocols to provide sufficient bandwidth and data integrity suitable for IT data network applications. Therefore, power line protocol designs have remained the province of proprietary applications and fully engineered control devices (e.g. X-10).

The IEEE P1901 effort, through the visibility that accompanies work of that standards body, may change the current proprietary nature of protocols using power lines as their medium. At present, the P1901 standards work is focused on the MAC, PHY and data link layers. The goal is to codify two PHY modulation techniques and allow their coexistence on the

same media-as contrasted to the multitude of individual IETF wireless standards that deal with protocols that share the same media (RF spectrum). At this point, the plans for the upper layers of the protocol are not clear and P1901 may occupy, in the power line domain, the same role as IEEE 802.15.4 does in the wireless domain. That is, a well-defined standard that codifies the lowest two or three layers of the protocol stack and can serve as the basis for both proprietary and open layers that can be incorporated into real world systems.

Table 5. OSI layers vs. Protocol Standards

OSI Network Layer	CeBus HomePlug 1.0 (Intellon)	HomePlug CC	IEEE P1901 HomePlug AV	Proprietary Insteon	UPB	Proprietary X-10 (Aztec)
Application	Application API included in standard	Application API included in standard	TBD	Application API included in standard	Application API included in standard	Application API included in standard
Presentation	Part of standard	Part of standard	TBD	Part of standard	Part of standard	Part of standard
Session	Part of standard	Part of standard	TBD	Part of standard	Part of standard	Part of standard
Transport	Part of standard	Part of standard	TBD	Part of standard	Part of standard	Part of standard
Network	Part of standard	Part of standard	TBD	Part of standard	Part of standard	Part of standard
DataLink	Part of standard	Part of standard	Hardware/software	Part of standard	Part of standard	Part of standard
Physical	Hardware	Hardware	Hardware	Hardware	Hardware	Hardware

Table 6. High-Level Protocol Characteristics

Protocol	Bandwidth	One/Two Way Communication	Security -Data Integrity -Protection from Data Loss -Confidentiality	Openness
HomePlug 1.0 (Intellon)	14 mbps/ 85 mbps (unofficial)	Two way	-considered very robust -security and data loss status unknown	Private, available to license holders
HomePlug CC	7.5kbps	Two way	-predicted to be very robust -security and data loss status unknown	Private, available to license holders
HomePlug AV (IEEE P1901)	>100mbps	Two way	Security characteristics expected similar to current IP protocols	Proposed open IEEE standard
Insteon	2.99kbps	One way	-marginal data integrity -no security or data loss protection	Private, at least one additional license holder
UPB	240bps	One way with ACK	-marginal data integrity -no security or data loss protection	Private, licensed to multiple manufacturers
-10	120 bps	One way and One way with ACK	-marginal data integrity -no security or data loss protection	Private

Table 7. Protocol Openness and Usability

Protocol	Protocol Specification Status	Protocol Status/ Completeness	Developer Protocol Interface	Deployment / Diagnostic Tools
HomePlug 1.0 (Intellon)	Written spec. available to licensees	Implementations shipping since 2002	MAC/PHY chip from Intellon	Yes
HomePlug CC	Written spec. available to licensees	Early chip level implementations shipping	MAC/PHY chip from Yitran	No
HomePlug AV (IEEE P1901)	Written submission to IEEE	Still in proposal stage of development	None at present, expect MAC/PHY chip level access	No
Insteon	Written spec.	Implementations shipping since 2004	TBD	Private, at least one additional license holder
UPB	Written spec. available to licensees	Implementations shipping since 2002, multiple manufacturers	Command interface to microcontroller	Yes, embedded diagnostic tools
X-10	Written spec.	One way and One way with ACK	-marginal data integrity -no security or data loss protection	Yes

Can We Integrate these Technologies?

In many ways, the vision is clear. Once enabled by a suitable control and data network, homes can be controlled to optimize convenience, comfort and security while they minimize energy use and meaningfully respond to utility load management requests. Present home networking technology, however, paints a very different picture. Communications technologies for currently installed residential lighting, heating/air conditioning and irrigation control systems have been selected with little regard for other technologies found in residential or IT systems. As a result, these separate, vertical applications operate side-by-side in relative isolation from each other, and that isolation extends throughout the system from sensors and actuators all the way up to user interfaces.

This pattern of isolated “islands of control” has been observed and, to some extent, addressed in the C & I building EMCS domain. In some large buildings, the installed equipment base is of sufficient value that it cannot simply be replaced. As a result, the additional engineering effort required to integrate existing systems into a new control system is justified and design accommodations are undertaken. There are two classic systems-level approaches followed when integrating disparate systems such as these into a single entity. One is to introduce a gateway that is capable of translating messages from each existing system into a common language and format. A new control system would then use this gateway to control existing systems using commands native to each separate system. While the gateway approach has been successful in a number of industrial settings, it is labor intensive and often fragile. It requires the central system to fully understand the command semantics and the communications protocols for all systems connected to the gateway and to control and operate them in a manner that may not be consistent with their original design intent. While this approach can be implemented, the engineering expense, mostly software effort, required to create a fully integrated system may be unacceptable.

A second approach also involves, as with the gateway model, the introduction of a new control system. But, rather than attempt the integration of existing systems in their entirety, the new system only uses existing actuators, sensors and, if possible, communications networks present in the existing systems – discarding the hardware controllers and behavioral models of the original system. While the new system makes efficient use of existing infrastructure, it is not burdened with the task of harmonizing multiple command and behavioral models for control and monitoring. And, in the case of implementing a building-wide demand response system, the new system’s response to a DR event can be implemented directly with no need to map or coerce commands onto the semantics of older systems that may have few, if any, DR capabilities.

What is the relevance of these commercial building scenarios to the domain of residential DR? The potential advantages to both utilities and residential customers suggest that residential DR systems will become an important part of the energy management landscape in the near future. The same justification for a single conceptual control model and communications protocol that has promoted OpenADR in the commercial world will argue strongly for the development of similar standard(s) in the residential DR domain. And, given the scaling issues present in the residential world, these efforts for a single market-unifying application protocol will be given further importance. The critical question is, once these newly standardized DR messages enter the residential envelope, what will they communicate with? At present, there are few home automation entities that are capable of being programmed to

respond to DR messages. And, of these devices, few of them have sufficient computational power to interpret and respond to an Internet-delivered XML message – the format used for OpenADR messaging.

Standardizing and promoting the residential computing and communications environment with which utility DRAS systems will communicate is the currently “missing element” of the residential DR picture. And, the lessons learned in creating a single, unified point of contact for commercial systems can be fruitfully applied to the residential domain. While some homes have modern, whole-house control systems, most have a collection of disjoint, isolated lighting and environmental automation/control systems. By unifying these systems through gateway technologies (e.g. a residential EMCS) similar to those found in commercial environments and presenting a standardized interface to the DR world, it will be possible to maximize the number of residences that can participate in evolving residential DR domain.

Conclusions and Future Research Issues

The current success of both demonstration and early commercialized C & I DR standardized communications systems has provided both the credibility and justification for pursuing similarly standardized DR solutions in the residential domain. Furthermore, the development of open application-level standards, such as OpenADR, has provided the conceptual framework and supporting use case analysis necessary for rapid progress in that domain. However, at present, the computing and control environment found in the home lacks much of the sophistication and capability present in the C & I domain. While this is due, historically, to the weak financial justification for costly residential EMCS capabilities, the situation is rapidly changing. Residential automation systems are being installed in increasing numbers and, due to technical advances in electronics and networking design, at reasonable costs. Many of these systems have sufficient functionality to act as cooperating clients in a distributed DR system. The critical “piece of the puzzle” that is currently lacking is a conceptual model or “reference design” that codifies the ways in which home automation systems interact with external DR systems. While it may remain necessary and, perhaps desirable, for internal home network technologies to vary from one household (or community) to another, the development of a common residential information model and the codification of how that information is passed between a home and a DR server will substantially benefit the widest group of utilities and households in the domain of residential demand response.

There are a number of pressing research questions in this domain. One issue is whether and, if so, how should the relatively rich semantics of OpenADR be mapped onto a simpler set of signals that can be presented to residential home automation gateways. Part of this issue is whether homes participating in DR programs should have the computational power to parse and interpret OpenADR XML messages. While the computational power required to parse these messages is only a fraction of that found in a typical PC, it must be available for OpenADR messages that could enter the home at any time and, so will require dedicated hardware. The computational needs for OpenADR clients may be substantially reduced if we can map the current event message content in a way that simplifies or eliminates complex message parsing and formatting. There are obvious ways to map these XML messages onto simpler representations. However, that mapping inevitably removes some functionality of the DR server↔client interface and some research effort would be useful in determining what level of functionality is appropriate or required.

As described above, there are a variety of networking technologies available within the residential envelope. Each has its own set of characteristics – speed, error rate, etc. A systematic study of the end use communications requirements should be undertaken to see which of these technologies are suitable for use in household control and monitoring applications. For example, what level of reliability is suitable for time-shifting dishwasher operations? A well-founded understanding of control requirements for these applications will drive the acceptance of various home networking technologies.

It is critical to understand the degree to which wireless and power line protocols can be used interchangeably and/or integrated together into a single “seamless” home automation network. A number of initiatives have begun in the area of integrated home automation control networking (e.g. OpenHAN) and the utility/aggregator↔home interface will be strongly influenced by these efforts. Residential systems represent a broad spectrum of functionality and, for both small, simple dwellings and large, multi-floor homes, the control models that describe them will enable (or limit) the creation of large DR systems. The development of a common, scalable information and control model for a wide range of US residential dwellings will, in no small way, guarantee the success of these efforts.

Acknowledgements

The work described in this report was funded by the Demand Response Research Center which is funded by the California Energy Commission (Energy Commission), Public Interest

Energy Research (PIER) Program, under Work for Others Contract No.500-03-026, Am #1 and by the U.S. Department of Energy under Contract No. DE-AC02-05CH11231.

References

[1] Koch, Piette “Architecture Concepts and Technical Issues for an Open, Interoperable Automated Demand Response Infrastructure”, Lawrence Berkeley National Laboratory Report LBNL-63664

[2] Intermediate Energy Infobook”, p. 46, The Need Project, 2008,
[http://www.need.org/needpdf/Intermediate Energy Infobook.pdf](http://www.need.org/needpdf/Intermediate_Energy_Infobook.pdf)

[3] Automated Demand Response Implementation Award, 2008 Buildy Award Winner at
<http://www.automatedbuildings.com/news/jun08/reviews/080528114043buildy.htm>

[4], [9] LBNL/DRRC OpenADR-2008-R2 downloaded at *<http://openadr.lbl.gov/index.html>*

[5] Han, Piette “Solutions for Summer Electric Power Shortages: Demand Response and its Applications in Air Conditioning and Refrigerating Systems”, Lawrence Berkeley National Laboratory report LBNL-63806

[6] *http://www.demandsidemanagement.com/energy_management_control_systems.htm*

[7] “Advancing the Efficiency of Electrical Utilization: ‘Prices to Devices’”, EPRI Energy Technology Assessment Center

[8] Residential Buildings Committee (Northern California Chapter) “Cripple Wall Fact Sheet”,
http://www.quake06/best_practices/fact_sheets/criple_wall.pdf

[10] “Internetworking Technology Handbook”,
<http://www.cisco.com/en/US/docs/internetworking/technologh/handbook>

[11] *http://www.e-principles.com/networks_in_the_home_the_new_growth_market.htm*

[12] National Radio Systems Committee “Radio Broadcast Data Standard”,
<http://www.nrscstandards.org/RBDS/default.asp>

[13] Building Industry Consulting Services International (BICSI) “Residential Network Cabling”. McGraw Hill ISBN 0071382119

[14] *<http://www.ieee802.org/11/index.shtml>*

[15] <http://www.ieee802.org/15/pub/TG1.html>, <http://www.ieee802.org/15/pub/TG2.html>,
<http://www.ieee802.org/15/pub/TG3.html>

[16] <http://www.ieee802.org/15/pub/TG4.html>

